Predicting Nearly As Well As the Optimal Twice Differentiable Regressor

N. Denizcan Vanli, Muhammed O. Sayin, and Suleyman S. Kozat, *Senior Member, IEEE*

arXiv:1401.6413v2 [cs.LG] 6 Oct 2014

***Abstract*—We study nonlinear regression of real valued data in an individual sequence manner, where we provide results that are guaranteed to hold without any statistical assumptions. We address the convergence and undertraining issues of conventional nonlinear regression methods and introduce an algorithm that elegantly mitigates these issues via an incremental hierarchical structure, (i.e., via an incremental decision tree). Particularly, we present a piecewise linear (or nonlinear) regression algorithm that partitions the regressor space in a data driven manner and learns a linear model at each region. Unlike the conventional approaches, our algorithm gradually increases the number of disjoint partitions on the regressor space in a sequential manner according to the observed data. Through this data driven approach, our algorithm sequentially and asymptotically achieves the performance of the optimal twice differentiable regression function for any data sequence with an unknown and arbitrary length. The computational complexity of the introduced algorithm is only logarithmic in the data length under certain regularity conditions. We provide the explicit description of the algorithm and demonstrate the significant gains for the well-known benchmark real data sets and chaotic signals.**

***Index Terms*—Online, nonlinear, regression, incremental decision tree.**

I. Introduction

We study sequential nonlinear regression, where we aim to estimate or model a desired sequence $\{d[t]\}_{t\geq 1}$ by using a sequence of regressor vectors $\{\boldsymbol{x}[t]\}_{t\geq 1}$. In particular, we seek to find the relationship, if it exists, between these two sequences, which is assumed to be unknown, nonlinear, and possibly time varying. This generic nonlinear regression framework is extensively studied in the machine learning and signal processing literatures since it can model a wide range of real life applications by capturing the salient characteristics of underlying signals and systems [1]–[15]. In order to define and find this relationship between the desired sequence and regressor vectors, numerous methods such as neural networks, Volterra filters, and B-splines are used [1], [2], [6], [11], [16]–[19]. However, either these methods are extremely difficult to use in real life applications due to convergence issues, e.g., Volterra filters and B-splines, or it is quite hard to obtain a consistent performance in different scenarios, cf. [1]–[3], [6]–[8], [12]–[15], [20]–[22].

To this end, in this paper, we propose an algorithm that alleviates these issues by introducing hierarchical models that recursively and effectively partition the regressor space into subsequent regions in a data driven manner, where a different linear model is learned at each region. Unlike most of the nonlinear models, learning linear structures at each region can be efficiently managed. Hence, using this hierarchical piecewise model, we significantly mitigate the convergence and consistency issues. Furthermore, we prove that the resulting hierarchical piecewise model asymptotically achieves the performance of any twice differentiable regression function that is directly tuned to the underlying observations without any tuning of algorithmic parameters or without any assumptions on the data (other than an upper bound on the magnitude). Since most of the nonlinear modeling functions of the regression algorithms in the literature, such as neural networks and Volterra filters, can be accurately represented by twice differentiable functions [1], [2], [6], [16], [18], [19], our algorithm readily performs asymptotically as well as such nonlinear learning algorithms.

In particular, the introduced method sequentially and recursively divides the space of the regressors into disjoint regions according to the amount of the data in each region, instead of committing to a priori selected partition. In this sense, we avoid creating undertrained regions until a sufficient amount of data is observed. The nonlinear modeling power of the introduced algorithm is incremented (by consecutively partitioning the regressor space into smaller regions) as the observed data length increases. The introduced method adapts itself according to the observed data instead of relying on ad-hoc parameters that are set while initializing the algorithm. Thus, the introduced algorithm provides a significantly stronger modeling power with respect to the state-of-the-art methods in the literature as shown in our experiments.

We emphasize that piecewise linear regression using tree structures is extensively studied in the computational learning and signal processing literatures [7]–[9], [12]–[15], [20]–[24] due to its attractive convergence and consistency features. There exist several tree based algorithms that mitigate the overtraining problem by defining hierarchical piecewise models such as [8], [9], [12]–[15]. Although these methods achieve the performance of the best piecewise model defined on a tree, i.e., the best pruning of a tree, they only yield satisfactory performance when the initial partitioning of the regressor space is highly accurate or tuned to the underlying data (which is unknown or even time-varying). Furthermore, there are more recent algorithms such as [20] that achieve the performance of the optimal combination of all piecewise models defined on a tree that minimizes the accumulated loss. There are also methods that alleviate the overtraining

This work was supported in part by the Turkish Academy of Sciences Outstanding Researcher Programme and in part by TUBITAK under Contract 112E161 and Contract 113E517.

The authors are with the Department of Electrical and Electronics Engineering, Bilkent University, Bilkent, Ankara 06800, Turkey (e-mail: {vanli,sayin,kozat}@ee.bilkent.edu.tr).

problem by learning the region boundaries [20] to minimize the regression error for a fixed depth tree with a computational complexity relatively greater compared to the ones in [8], [9], [12], [14], [15] (particularly, exponential in the depth of the tree). However, these algorithms can only provide a limited modeling power since the tree structure in these studies is fixed. Furthermore, the methods such as [20] can only learn the locally optimal region boundaries due to the highly nonlinear (and non-convex) optimization structure. Unlike these methods, the introduced algorithm sequentially increases its nonlinear modeling power according to the observed data and directly achieves the performance of the best twice differentiable regression function that minimizes the accumulated regression error. We also show that in order to achieve the performance of a finer piecewise model defined on a tree, it is not even necessary to create these piecewise models when initializing the algorithm. Hence, we do not train a piecewise model until a sufficient amount of data is observed, and show that the introduced algorithm, in this manner, does not suffer any asymptotical performance degradation. Therefore, unlike the relevant studies in the literature, in which undertrained (i.e., unnecessary) partitions are kept in the overall structure, our method intrinsically eliminates the unnecessarily finer partitions without any loss in asymptotical performance (i.e., we maintain universality).

Aside from such piecewise linear regression techniques based on hierarchical models, there are various different methods to introduce nonlinearity such as B-splines and Volterra series [1], [2], [6], [11], [17]–[19]. In these methods, the nonlinearity is usually introduced by modifying the basis functions to create polynomial estimators, e.g., in [19], the authors use trigonometric functions as their basis functions. We emphasize that these techniques can be straightforwardly incorporated into our framework by using these methods at each region in the introduced algorithm to obtain piecewise nonlinear regressors. Note that the performance of such methods, e.g., B-splines and Volterra series (and other various methods with different basis functions), is satisfactory when the data is generated using the underlying basis functions of the regressor. In real life applications, the underlying model that generates the data is usually unknown. Thus, the successful implementation of these methods significantly depends on the match (or mismatch) between the regressor structure and the underlying model generating the data. On the other hand, the introduced algorithm achieves the performance of any such regressor provided that its basis functions are twice differentiable. In this sense, unlike the conventional methods in the literature, whose performances are highly dependant on the selection of the basis functions, our method can well approximate these basis functions (and regressors formed by these basis functions) via piecewise models such that the performance difference with respect to the best such regressor asymptotically goes to zero in a strong individual sequence manner without any statistical assumptions.

The main contributions of this paper are as follows. We introduce a sequential piecewise linear regression algorithm *i)* that provides a significantly improved modeling power by adaptively increasing the number of partitions according to the observed data, *ii)* that is highly efficient in terms of the computational complexity as well as the error performance, and *iii)* whose performance converges to *iii-a)* the performance of the optimal twice differentiable function that is selected in hindsight and *iii-b)* the best piecewise linear model defined on the incremental decision tree, with guaranteed upper bounds without any statistical or structural assumptions on the desired data as well as on the regressor vectors (other than an upper bound on them). Hence, unlike the state-of-the-art approaches whose performances usually depend on the initial construction of the tree, we introduce a method to construct a decision tree, whose depth (and structure) is adaptively incremented (and adjusted) in a data dependent manner, which we call an incremental decision tree. Furthermore, the introduced algorithm achieves this superior performance only with a computational complexity $O(\log(n))$ for any data length n, under certain regularity conditions. Even if these regularity conditions are not met, the introduced algorithm still achieves the performance of any twice differentiable regression function, however with a computational complexity linear in the data length.

The organization of the paper is as follows. We first describe the sequential piecewise linear regression problem in detail in Section II. We then introduce the main algorithm in Section III and prove that the performance of this algorithm is nearly as well as the best piecewise linear model that can be defined by the incremental decision tree in Section IV. Using this result, we also show that the introduced algorithm achieves the performance of the optimal twice differentiable function that is selected after observing the entire data before processing starts, i.e., non-causally. In Section V, we demonstrate the performance of the introduced algorithm through simulations and then conclude the paper with several remarks in Section VI.

II. Problem Description

We study sequential nonlinear regression, where the aim is to estimate an unknown desired sequence $\{d[t]\}_{t\geq 1}$ by using a sequence of regressor vectors $\{\boldsymbol{x}[t]\}_{t\geq 1}$, where the desired sequence and the regressor vectors are real valued and bounded but otherwise arbitrary, i.e., $d[t] \in \mathbb{R}$, $\boldsymbol{x}[t] \triangleq [x_1[t],\ldots,x_p[t]]^T \in \mathbb{R}^p$ for an arbitrary integer p and $|d[t]|, |x_i[t]| < A < \infty$ for all t and $i = 1,\ldots,p$. We call the regressors as "sequential" if in order to estimate the desired data at time t, i.e., $d[t]$, they only use the past information $d[1],\ldots,d[t-1]$ and the observed regressor vectors[1] $\boldsymbol{x}[1],\ldots,\boldsymbol{x}[t]$.

In this framework, a piecewise linear model is constructed by dividing the regressor space into a union of disjoint regions, where in each region a linear model holds. As an example, suppose that the regressor space is parsed into K disjoint regions $\mathcal{R}_1,\ldots,\mathcal{R}_K$ such that $\bigcup_{k=1}^{K}\mathcal{R}_k = [-A,A]^p$. Given such a model, say model m, at each time t, the sequential

[1]All vectors are column vectors and denoted by boldface lower case letters. Matrices are denoted by boldface upper case letters. For a vector $\boldsymbol{x}$, $\boldsymbol{x}^T$ is the ordinary transpose. We denote $d_a^b \triangleq \{d[t]\}_{t=a}^b$.

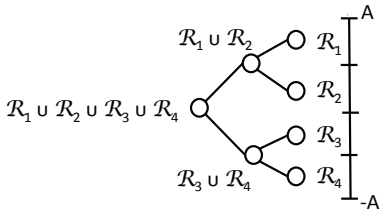

Fig. 1: The partitioning of a one dimensional regressor space, i.e., $[-A, A]$, using a depth-2 full decision tree, where each node represents a portion of the regressor space.

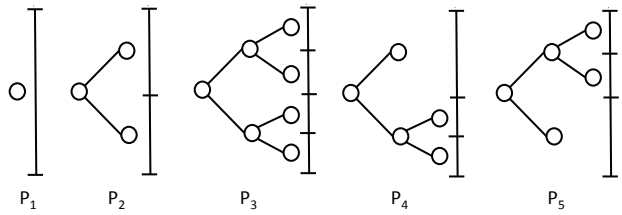

Fig. 2: All different piecewise linear models that can be obtained using a depth-2 full decision tree, where the regressor space is one dimensional. These models are based on the partitioning shown in Fig. 1.

linear[2] regressor predicts $d[t]$ as $\hat{d}_m[t] = \boldsymbol{v}_{m,k}^T[t]\boldsymbol{x}[t]$ when $\boldsymbol{x}[t] \in \mathcal{R}_k$, where $\boldsymbol{v}_{m,k}[t] \in \mathbb{R}^p$ for all $k = 1, \ldots, K$. These linear models assigned to each region can be trained independently using different adaptive methods such as the least mean squares (LMS) or the recursive least squares (RLS) algorithms.

However, by directly partitioning the regressor space as $\bigcup_{k=1}^{K} \mathcal{R}_k = [-A, A]^p$ before the processing starts and optimizing only the internal parameters of the piecewise linear model, i.e., $\boldsymbol{v}_{m,k}[t]$, one significantly limits the performance of the overall regressor since we do not have any prior knowledge on the underlying desired signal. Therefore, instead of committing to a single piecewise linear model with a fixed and given partitioning, and performing optimization only over the internal linear regression parameters of this regressor, one can use a decision tree to partition the regressor space and try to achieve the performance of the best partitioning over the whole doubly exponential number of different models represented by this tree [25].

As an example, in Fig. 1, we partition the one dimensional regressor space $[-A, A]$, using a depth-2 tree, where the regions $\mathcal{R}_1, \ldots, \mathcal{R}_4$ correspond to disjoint intervals on the real line and the internal nodes are constructed using union of these regions. In the generic case, for a depth-d full decision tree, there exist 2^d leaf nodes and $2^d - 1$ internal nodes. Each node of the tree represents a portion of the regressor space such that the union of the regions represented by the leaf nodes is equal to the entire regressor space $[-A, A]^p$. Moreover, the region corresponding to each internal node is constructed by the union of the regions of its children. In this sense, we obtain $2^{d+1} - 1$ different nodes (regions) on the depth-d decision tree (on the regressor space) and approximately 1.5^{2^d} different piecewise models that can be represented by certain collections of the regions represented by the nodes of the decision tree [25]. For example, we consider the same scenario as in Fig. 1, where we partition the one dimensional real space using a depth-2 tree. Then, as shown in Fig. 1, there are 7 different nodes on the depth-2 decision tree; and as shown in Fig. 2, a depth-2 tree defines 5 different piecewise partitions or models, where each of these models is constructed using certain unions of the nodes of the full depth decision tree.

We emphasize that given a decision tree of depth-d, the nonlinear modeling power of this tree is fixed and finite since there are only $2^{d+1} - 1$ different regions (one for each node) and approximately 1.5^{2^d} different piecewise models (i.e., partitions) defined on this tree. Instead of introducing such a limitation, we recursively increment the depth of the decision tree as the data length increases. We call such a tree the "incremental decision tree" since the depth of the decision tree is incremented (and potentially goes to infinity) as the data length n increases, hence in a certain sense, we can achieve the modeling power of an infinite depth tree. As shown in Theorem 2, the piecewise linear models defined on the tree will converge to any unknown underlying twice differentiable model under certain regularity conditions as n increases.

To this end, we seek to find a sequential regression algorithm (whose estimate at time t is represented by $\hat{d}_s[t]$), when applied to any sequence of data and regressor vectors, yields the following performance (i.e., regret) guarantee

$$\sum_{t=1}^{n} \left(d[t] - \hat{d}_s[t]\right)^2 - \inf_{f \in \mathcal{F}} \sum_{t=1}^{n} \left(d[t] - \hat{d}_f[t]\right)^2 \leq o(n), \quad (1)$$

over any n, without the knowledge of n, where $\mathcal{F}$ represents the class of all twice differentiable functions, whose parameters are set in hindsight, i.e., after observing the entire data before processing starts, and $\hat{d}_f[t]$ represents the estimate of the twice differentiable function $f \in \mathcal{F}$ at time t. The relative accumulated error in (1) represents the performance difference of the introduced algorithm and the optimal batch twice differentiable regressor. Hence, an upper bound of $o(n)$ in (1) implies that the algorithm $\hat{d}_s[t]$ sequentially and asymptotically converges to the performance of the regressor $\hat{d}_f[t]$, for any $f \in \mathcal{F}$.

III. Nonlinear Regression via Incremental Decision Trees

In this section, we introduce the main results of the paper. Particularly, we first show that the introduced sequential piecewise linear regression algorithm asymptotically achieves the performance of the best piecewise linear model defined on the incremental decision tree (with possibly infinite depth) with the optimal regression parameters at each region that minimizes the accumulated loss. We then use this result to prove that the introduced algorithm asymptotically achieves the performance of any twice differentiable regression function. We provide the algorithmic details and the construction of the algorithm in Section IV.

[2]Note that affine models can also be represented as linear models by appending a 1 to $\boldsymbol{x}[t]$, where the dimension of the regressor space increases by one.

Theorem 1: *Let $\{d[t]\}_{t\geq 1}$ and $\{\boldsymbol{x}[t]\}_{t\geq 1}$ be arbitrary, bounded, and real-valued sequences of data and regressor vectors, respectively. Then the algorithm $\hat{d}[t]$ (given in Fig. 4) when applied to these data sequences yields*

$$\sum_{t=1}^{n}\left(d[t]-\hat{d}[t]\right)^2 - \inf_{m\in\mathcal{M}'_n}\left[\inf_{\substack{\boldsymbol{v}_{m,k}\in\mathbb{R}^p\\ k=1,\ldots,K_m}}\left\{\sum_{t=1}^{n}\left(d[t]-\hat{d}_b[t]\right)^2+\delta\,||\boldsymbol{v}_m||^2\right\}\right] \leq O\left(p\log^2(n)\right),$$

for any n, with a computational complexity upper bounded by $O(n)$, where K_m denotes the number of leaf nodes in the hierarchical model m, $\mathcal{M}_n$ represents the set of all hierarchical models defined on the incremental decision tree at time n, $\mathcal{M}'_n$ represents the set of all hierarchical models with at most $O(\log(n))$ leaves defined on the incremental decision tree at time n, i.e., $\mathcal{M}'_n \triangleq \{m\in\mathcal{M}_n : K_m \leq O(\log(n))\}$, and $\boldsymbol{v}_m \triangleq [\boldsymbol{v}_{m,1};\ldots;\boldsymbol{v}_{m,K_m}]$.

This theorem indicates that the introduced algorithm can asymptotically and sequentially achieve the performance of any piecewise model in the set $\mathcal{M}'_n$, i.e., the piecewise models having at most $O(\log(n))$ leaves defined on the tree. In particular, over any unknown data length n, the performance of the piecewise models with $O(\log(n))$ leaves can be sequentially achieved by the introduced algorithm with a regret upper bounded by $O\left(p\log^2(n)\right)$. In this sense, we do not compare the performance of the introduced algorithm with a fixed class of regressors, over any data length n. Instead, the regret of the introduced algorithm is defined with respect to a set of piecewise linear regressors, whose number of partitions are upper bounded by $O(\log(n))$, i.e., the competition class grows as n increases. In the conventional tree based regression methods, the depth of the tree is set before processing starts and the performance of the regressor is highly sensitive with respect to the unknown data length. For example, if the depth of the tree is large whereas there are not enough data samples, then the piecewise model will be undertrained and yield an unsatisfactory performance. Similarly, if the depth of the tree is small whereas huge number of data samples are available, then trees (and regressors) with higher depths (and finer regions) can be better trained. As shown in Theorem 1, the introduced algorithm elegantly and intrinsically makes such decisions and performs asymptotically as well as any piecewise regressor in the competition class that grows exponentially in n [25]. Such a significant performance is achieved with a computational complexity upper bounded by $O(n)$, i.e., only linear in the data length, whereas the number of different piecewise models defined on the incremental decision tree can be in the order of 1.5^n [25]. Moreover, under certain regularity conditions the computational complexity of the algorithm is $O(\log(n))$ as will be discussed in Remark 2. This theorem is an intermediate step to show that the introduced algorithm yields the desired performance guarantee in (1), and will be used to prove the next theorem.

Using Theorem 1, we introduce another theorem presenting the main result of the paper, where we define the performance of the introduced algorithm with respect to the class of twice differentiable functions as in (1).

Theorem 2: *Let $\{d[t]\}_{t\geq 1}$ and $\{\boldsymbol{x}[t]\}_{t\geq 1}$ be arbitrary, bounded, and real-valued sequences of data and regressor vectors, respectively. Let $\mathcal{F}$ be the class of all twice differentiable functions such that for any $f\in\mathcal{F}$, $\frac{\partial^2 f(\boldsymbol{x})}{\partial x_i \partial x_j}\leq D<\infty$, $i,j=1,\ldots,p$ and we denote $\hat{d}_f[t]=f(\boldsymbol{x}[t])$. Then the algorithm $\hat{d}[t]$ given in Fig. 4 when applied to these data sequences yields*

$$\sum_{t=1}^{n}\left(d[t]-\hat{d}[t]\right)^2 - \inf_{f\in\mathcal{F}}\sum_{t=1}^{n}\left(d[t]-\hat{d}_f[t]\right)^2 \leq o(p^2 n),$$

for any n, with a computational complexity upper bounded by $O(n)$.

This theorem presents the nonlinear modeling power of the introduced algorithm. Specifically, it states that the introduced algorithm can asymptotically achieve the performance of the optimal twice differentiable function that is selected after observing the entire data in hindsight. Note that there are several kernel and neural network based sequential nonlinear regression algorithms [1], [2], [6] (which can be modeled via twice differentiable functions) whose computational complexities are similar to the introduced algorithm. However, the performances of such nonlinear models are only comparable with respect to their batch variants. On the other hand, we demonstrate the performance of the introduced algorithm with respect to a extremely large class of regressors without any statistical assumptions. In this sense, the performance of any regression algorithm that can be modeled by twice differentiable functions is asymptotically achievable by the introduced algorithm. Hence, the introduced algorithm yields a significantly more robust performance with respect to the such conventional approaches in the literature as also illustrated in different experiments in Section V.

The proofs of Theorem 1, Theorem 2, and the construction of the algorithm are given in the following section.

IV. Construction of the Algorithm and Proofs of the Theorems

In this section, we first introduce a labeling to efficiently manage the hierarchical models and then describe the algorithm in its main lines. We next prove Theorem 1, where we also provide the complete construction of the algorithm. We then present a proof for Theorem 2, using the results of Theorem 1.

A. Notation

We first introduce a labeling for the tree nodes following [26]. The root node is labeled with an empty binary string λ and assuming that a node has a label κ, where $\kappa=\nu_1\ldots\nu_l$ is a binary string of length l formed from letters $\nu_1,\ldots,\nu_l$, we label its upper and lower children as $\kappa 1$ and $\kappa 0$, respectively. Here, we emphasize that a string can only take its letters from the binary alphabet, i.e., $\nu\in\{0,1\}$, where 0 refers to the

lower child, and 1 refers to the upper child of a node. We also introduce another concept, i.e., the definition of the prefix of a string. We say that a string $\kappa' = \nu'_1 \dots \nu'_{l'}$ is a prefix to string $\kappa = \nu_1 \dots \nu_l$ if $l' \leq l$ and $\nu'_i = \nu_i$ for all $i = 1, \dots, l'$, and the empty string λ is a prefix to all strings. Finally, we let $\mathcal{P}(\kappa)$ represent all prefixes to the string κ, i.e., $\mathcal{P}(\kappa) \triangleq \{\kappa_0, \dots, \kappa_l\}$, where $l \triangleq l(\kappa)$ is the length of the string κ, κ_i is the string with $l(\kappa_i) = i$, and $\kappa_0 = \lambda$ is the empty string, such that the first i letters of the string κ forms the string κ_i for $i = 0, \dots, l$. Letting $\mathcal{L}$ denote the set of leaf nodes for a given decision tree, each leaf node of the tree, i.e., $\kappa \in \mathcal{L}$, is given a specific index $\alpha_\kappa \in \{0, \dots, M-1\}$ representing the number of regressor vectors that has fallen into $\mathcal{R}_\kappa$. For presentation purposes, we consider $M = 2$ throughout the paper.

B. Outline of the Algorithm

At time $t = 0$, the introduced algorithm starts with a single node (i.e., the root node) representing the entire regressor space. As the new data is observed, the proposed algorithm sequentially divides the regressor space into smaller disjoint regions according to the observed regressor vectors. In particular, each region is divided into subsequent child regions as soon as a new regressor vector has fallen into that region. In this incremental hierarchical structure, we assign an independent linear regressor to each node (i.e., to each region). Such a hierarchical structure (embedded with linear regressors) can define 1.5^n different piecewise linear models or partitions. We then combine the outputs of all these different piecewise models via a mixture of experts approach instead of committing to a single model. However, even for a small n, the number of piecewise models (i.e., experts) grows extremely rapidly (particularly, exponential in n). Hence, in order to perform this calculation in an efficient manner, we assign a weight to each node on the tree and present a method to calculate the final output using these weights with a significantly reduced computational complexity, i.e., logarithmic in n under certain regularity conditions.

We then compare the performance of the introduced algorithm with respect to the best batch piecewise model defined on the incremental decision tree. Our algorithm first suffers a "constructional regret" that arise from the adaptive construction of the incremental decision tree (since the finer piecewise models are not present at the beginning of the processing) and from the sequential combination of the outputs of all piecewise models (i.e., due to the mixture of experts approach). Second, each piecewise model suffers a "parameter regret" while sequentially learning the true regression parameters at each region. We provide deterministic upper bounds on these regrets and illustrate that the introduced algorithm is twice-universal, i.e., universal in both entire piecewise models (even though the finer models appear as n increases and do not used until then) and linear regression parameters.

C. Proof of Theorem 1 and Construction of the Algorithm

In this section, we describe the algorithm in detail and derive a regret upper bound with respect to the best batch piecewise model defined on the incremental decision tree.

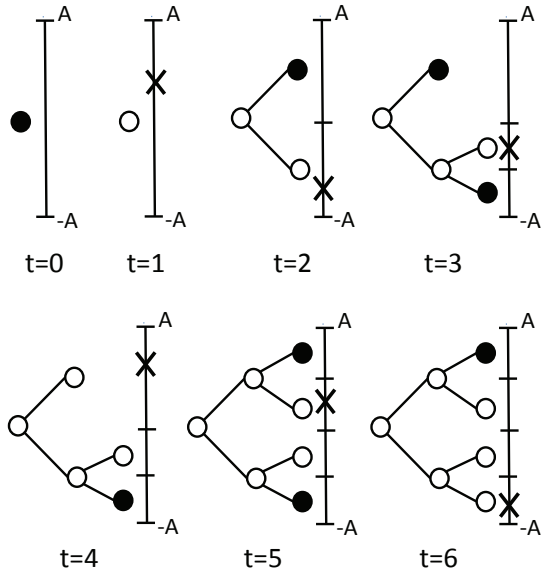

Fig. 3: A sample evolution of the incremental decision tree, where the regressor space is one dimensional. The "$\times$" marks on the regressor space represents the value of the regressor vector at that specific time instant. Light nodes are the ones having an index of 1, whereas the index of the dark nodes is 0.

Before the processing starts, i.e., at time $t = 0$, we begin with a single node, i.e., the root node λ, having index $\alpha_\lambda = 0$. Then, we recursively construct the decision tree according to the following principle. For every time instant $t > 0$, we find the leaf node of the tree $\kappa \in \mathcal{L}$ such that $\boldsymbol{x}[t] \in \mathcal{R}_\kappa$. For this node, if we have $\alpha_\kappa = 0$, we do not modify the tree but only increment this index by 1. On the other hand, if $\alpha_\kappa = 1$, then we generate two children nodes $\kappa 0, \kappa 1$ for this node by dividing the region $\mathcal{R}_\kappa$ into two disjoint regions $\mathcal{R}_{\kappa 0}, \mathcal{R}_{\kappa 1}$, using the plane $x_i = c$, where $i - 1 \equiv l(\kappa) \pmod p$ and c is the midpoint of the region $\mathcal{R}_\kappa$ along the ith dimension. For node $\kappa\nu$ with $\boldsymbol{x}[t] \in \mathcal{R}_{\kappa\nu}$ (i.e., the children node containing the current regressor vector), we set $\alpha_{\kappa\nu} = 1$ and the index of the other child is set to 0. The accumulated regressor vectors and the data in node κ are also transferred to its children to train a linear regressor in these child nodes.

As an example, in Fig. 3, we consider that the regressor space is one dimensional, i.e., $[-A, A]$, and present a sample evolution of the tree. In the figure, the nodes having an index of 0 are shown as dark nodes, whereas the others are light nodes, and the regressor vectors are marked with $\times$'s in the one dimensional regressor space. For instance at time $t = 2$, we have a depth-1 tree, where we have two nodes 0 and 1 with corresponding regions $\mathcal{R}_0 = [-A, 0]$, $\mathcal{R}_1 = [0, A]$, and $\alpha_0 = 1$, $\alpha_1 = 0$. Then, at time $t = 3$, we observe a regressor vector $\boldsymbol{x}[3] \in \mathcal{R}_0$ and divide this region into two disjoint regions using $x_1 = -A/2$ line. We then find that $\boldsymbol{x}[3] \in \mathcal{R}_{01}$, hence set $\alpha_{01} = 1$, whereas $\alpha_{00} = 0$.

We assign an independent linear regressor to each node on the incremental decision tree. Each linear regressor is trained using only the information contained in its corresponding node. Hence, we can obtain different piecewise models by using a certain collection of this node regressors according to the hierarchical structure. Each such piecewise model suffers

a parameter regret in order to sequentially learn the optimal linear regression parameters at each region that minimizes the cumulative error. This issue is discussed towards the end of this section.

Using this incremental hierarchical structure with linear regressors at each region, the incremental decision tree can represent up to 1.5^n different piecewise linear models after observing data of length n. For example, in Fig. 3, at time $t = 6$, we have 5 different piecewise linear models (see Fig. 2), whereas at time $t = 4$, we have 3 different piecewise linear models. Each of these piecewise linear models can be used to perform the estimation task. However, we use a mixture of experts approach to combine the outputs of all piecewise linear models, instead of choosing a single one among them.

To this end, one can assign a performance dependent weight to each piecewise linear model defined on the incremental decision tree and combine their weighted outputs to obtain the final estimate of the algorithm [16], [27], [28]. In a conventional setting, such a mixture of expert approach is guaranteed to asymptotically achieve the performance of the best piecewise linear model defined on the tree [16], [27], [28]. However, in our framework, to achieve the performance of the best twice differentiable regression function, as t increases (i.e., we observe new data), the total number of different piecewise linear models can increase exponentially with t. In this sense, we have a highly dynamic optimization framework. For example, in Fig. 3, at time $t = 4$, we have 3 different piecewise linear models, hence calculate the final output of our algorithm as $\hat{d}[t] = w_1[t]\hat{d}_1[t] + w_2[t]\hat{d}_2[t] + w_2[t]\hat{d}_2[t]$, where $\hat{d}_i[t]$ represents the output of the ith piecewise linear model and $w_i[t]$ represents its weight. However, at time $t = 6$, we have 5 different piecewise linear models, i.e., $\hat{d}[t] = \sum_{i=1}^{5} w_i[t]\hat{d}_i[t]$, therefore the number of experts increases. Hence, not only such a combination approach requires the processing of the entire observed data at each time t (i.e., it results in a brute-force batch-to-online conversion), but also it cannot be practically implemented even for a considerably short data sequences such as $n = 100$.

To elegantly solve this problem, we assign a weight to each node on the incremental decision tree, instead of using a conventional mixture of experts approach. In this way, we illustrate a method to calculate the original highly dynamic combination weights in an efficient manner, i.e., without requiring the processing of the entire data for each new sample, and with a significantly reduced computational complexity.

To accomplish this, to each leaf node $\kappa \in \mathcal{L}$, we assign a performance dependant weight [26] as follows

$$P_\kappa(n) \triangleq \exp\left\{-\frac{1}{2a}\sum_{t\leq n\,:\,\boldsymbol{x}[t]\in\mathcal{R}_\kappa}\left(d[t]-\hat{d}_{m,k}[t]\right)^2\right\},$$

where $\hat{d}_{m,k}[t]$ represents the linear regressor assigned to the kth node of the mth piecewise model and is constructed using the regressor introduced in [29] and discussed in (7). Then, we define the weight of an inner node $\kappa \notin \mathcal{L}$ as follows [26]

$$P_\kappa(n) \triangleq \frac{1}{2}P_{\kappa 0}(n)P_{\kappa 1}(n) + \frac{1}{2}\exp\left\{-\frac{1}{2a}\sum_{t\leq n\,:\,\boldsymbol{x}[t]\in\mathcal{R}_\kappa}\left(d[t]-\hat{d}_{m,k}[t]\right)^2\right\}.$$

Using this definitions, the weight of the root node λ can be constructed as follows

$$P_\lambda(n) = \sum_{m\in\mathcal{M}_n} 2^{-B_m}P(n|m),$$

where

$$P(n|m) \triangleq \exp\left\{-\frac{1}{2a}\sum_{t=1}^{n}\left(d[t]-\hat{d}_m[t]\right)^2\right\}$$

represents the performance of a given partition $m \in \mathcal{M}_n$ over a data length of n, and B_m represents the number of bits required to represent the model m on the binary tree using a universal code [30].

Hence, the performance of the root node satisfies $P_\lambda(n) \geq 2^{-B_m}P(n|m)$ for any $m \in \mathcal{M}_n$. That is,

$$-2a\ln\left(P_\lambda(n)\right) \leq \min_{m\in\mathcal{M}_n}\left\{\sum_{t=1}^{n}\left(d[t]-\hat{d}_m[t]\right)^2\right\} + 2a\ln(2)\log(n) + 4A^2K_m\log(n), \quad (2)$$

where the last line follows when we maximize B_m with respect to $m \in \mathcal{M}_n$ and the regret term $4A^2K_m\log(n)$ follows due to the adaptive construction of the incremental decision tree. This upper bound corresponds to the constructional regret of our algorithm.

Hence, we have obtained a weighting assignment achieving the performance of the optimal piecewise linear model. We next introduce a sequential algorithm achieving $P_\lambda(n)$. To this end, we first note that we have

$$P_\lambda(n) = \prod_{t=1}^{n}\frac{P_\lambda(t)}{P_\lambda(t-1)}. \quad (3)$$

Now if we can demonstrate a sequential algorithm whose performance is greater than or equal to $P_\lambda(t)/P_\lambda(t-1)$ for all t, we can conclude the proof. To this end, we present a sequential update from $P_\lambda(t-1)$ to $P_\lambda(t)$.

After the structural updates, i.e., the growth of the incremental decision tree, are completed, say at time t, we observe a regressor vector $\boldsymbol{x}[t] \in \mathcal{R}_\kappa$ for some $\kappa \in \mathcal{L}$. Then, we can compactly denote the weight of the root node at time $t-1$ as follows

$$P_\lambda(t-1) = \sum_{\kappa_i\in\mathcal{P}(\kappa)}\pi_{\kappa_i}[t-1]\exp\left\{-\frac{1}{2a}\sum_{\substack{t'<t\\ \boldsymbol{x}[t']\in\mathcal{R}_{\kappa_i}}}\left(d[t']-\hat{d}_{\kappa_i}[t']\right)^2\right\},$$

where $\hat{d}_\kappa[t]$ represents the output of the regressor for node κ, $\kappa_i \in \mathcal{P}(\kappa)$ is the string formed from the first i letters of $\kappa = \nu_1 \ldots \nu_l$, and $\pi_{\kappa_i}[t]$ is recursively defined as follows

$$\pi_{\kappa_i}[t] \triangleq \begin{cases}\frac{1}{2} & \text{, if } i = 0\\ \frac{1}{2}P_{\kappa_{i-1}\nu_i^c}(t-1)\pi_{\kappa_{i-1}}[t] & \text{, if } 1 \leq i \leq l-1\\ P_{\kappa_{i-1}\nu_i^c}(t-1)\pi_{\kappa_{i-1}}[t] & \text{, if } i = l\end{cases}.$$

Since $\boldsymbol{x}[t] \in \mathcal{R}_\kappa$ for some $\kappa \in \mathcal{L}$, then after $d[t]$ is revealed, the weight of the root node at time t can be calculated as follows

$$P_\lambda(t) = \sum_{\kappa_i \in \mathcal{P}(\kappa)} \pi_{\kappa_i}[t-1] \exp\left\{-\frac{1}{2a}\left(d[t]-\hat{d}_{\kappa_i}[t]\right)^2\right\} \times \exp\left\{-\frac{1}{2a}\sum_{t'<t\,:\,\boldsymbol{x}[t']\in\mathcal{R}_{\kappa_i}}\left(d[t']-\hat{d}_{\kappa_i}[t']\right)^2\right\},$$

which results in

$$\frac{P_\lambda(t)}{P_\lambda(t-1)} = \sum_{\kappa_i \in \mathcal{P}(\kappa)} \mu_{\kappa_i}[t-1] \exp\left\{-\frac{1}{2a}\left(d[t]-\hat{d}_{\kappa_i}[t]\right)^2\right\}, \tag{4}$$

where

$$\mu_{\kappa_i}[t-1] \triangleq \frac{\pi_{\kappa_i}[t-1] \exp\left\{-\frac{1}{2a}\sum_{\boldsymbol{x}[t']\in\mathcal{R}_{\kappa_i}}\left(d[t']-\hat{d}_{\kappa_i}[t']\right)^2\right\}}{P_\lambda(t-1)}.$$

We then focus on (4) and observe that we have $\sum_{\kappa_i \in \mathcal{P}(\kappa)} \mu_{\kappa_i}[t-1] = 1$, which means that if the second term in (4), i.e.,

$$f(\hat{d}_{\kappa_i}[t]) \triangleq \exp\left\{-\frac{1}{2a}\left(d[t]-\hat{d}_{\kappa_i}[t]\right)^2\right\},$$

is concave, then by Jensen's inequality, we can conclude that

$$\exp\left\{-\frac{1}{2a}\left(d[t]-\sum_{\kappa_i \in \mathcal{P}(\kappa)} \mu_{\kappa_i}[t-1]\hat{d}_{\kappa_i}[t]\right)^2\right\} \geq P_\lambda(t\,|\,t-1). \tag{5}$$

Since the function $f(\hat{d}_{\kappa_i}[t])$ is concave when $\left(d[t]-\hat{d}_{\kappa_i}[t]\right)^2 < a$, and we have $|d[t]| \leq A$, we have to set $a \geq 4A^2$. Therefore, we obtain a sequential regressor in (5), whose performance is greater than or equal to the performance of the root node, and the final estimate of our algorithm is calculated as follows

$$\hat{d}[t] \triangleq \sum_{\kappa_i \in \mathcal{P}(\kappa)} \mu_{\kappa_i}[t-1]\hat{d}_{\kappa_i}[t]. \tag{6}$$

Hence, our algorithm can achieve the performance of the best piecewise linear model defined on the incremental tree with a constructional regret given in (2). In order to achieve the performance of the best "batch" piecewise linear model, the introduced algorithm also suffers a parameter regret while learning the true regression parameters at each region. An upper bound on this regret is calculated as follows.

Consider an arbitrary piecewise model defined on the incremental decision tree, say the mth model, having K_m disjoint regions $\mathcal{R}_1, \ldots, \mathcal{R}_{K_m}$ such that $\bigcup_{k=1}^{K_m} \mathcal{R}_k = [-A, A]^p$. Then, a piecewise linear regressor can be constructed using the universal linear predictor of [29] in each region as $\hat{d}_m[t] = \boldsymbol{v}_{m,k}^T[t]\,\boldsymbol{x}[t]$, when $\boldsymbol{x}[t] \in \mathcal{R}_k$, with the regression parameters $\boldsymbol{v}_{m,k}[t] = \left(\boldsymbol{R}_k[t] + \delta \boldsymbol{I}\right)^{-1} \boldsymbol{p}_k[t]$, where $\boldsymbol{I}$ represents the appropriate sized identity matrix, $\boldsymbol{R}_k[t] \triangleq \sum_{t' \leq t\,:\,\boldsymbol{x}[t']\in\mathcal{R}_k} \boldsymbol{x}[t']\,\boldsymbol{x}^T[t']$, and $\boldsymbol{p}_k[t] \triangleq \sum_{t'<t\,:\,\boldsymbol{x}[t']\in\mathcal{R}_k} d[t']\,\boldsymbol{x}[t']$. The upper bound on the performance of this regressor can be calculated following similar lines to [29] and it is obtained as follows

$$\sum_{t=1}^{n}\left(d[t]-\hat{d}_m[t]\right)^2 - \min_{\substack{\boldsymbol{v}_{m,k}\in\mathbb{R}^p \\ k=1,\ldots,K_m}} \left\{\sum_{t=1}^{n}\left(d[t]-\hat{d}_b[t]\right)^2 + \delta\,||\boldsymbol{v}_m||^2\right\} \leq A^2 K_m p \ln(n/K_m) + O(1). \tag{7}$$

We emphasize that in each region of a piecewise model, different learning algorithms, e.g., different linear regressors or nonlinear ones, from the broad literature can be used. Note that although the main contribution of the paper is the hierarchical organization and efficient management of these piecewise models, we also discuss the implementation of a piecewise linear model [29] into our framework for completeness.

Finally, we achieve an upper bound on the performance of the introduced algorithm with respect to the best batch piecewise linear model. Combining the results in (2) and (7), we obtain

$$\begin{aligned}\sum_{t=1}^{n}\left(d[t]-\hat{d}[t]\right)^2 &\leq \min_{m\in\mathcal{M}_n}\left\{\sum_{t=1}^{n}\left(d[t]-\hat{d}_m[t]\right)^2\right\} + 2a\ln(2)\log(n) + 4A^2K_m\log(n) \\ &\leq \min_{m\in\mathcal{M}_n}\left[\min_{\substack{\boldsymbol{v}_{m,k}\in\mathbb{R}^p \\ k=1,\ldots,K_m}}\left\{\sum_{t=1}^{n}\left(d[t]-\hat{d}_b[t]\right)^2 + \delta\,||\boldsymbol{v}_m||^2\right\}\right] \\ &\quad + \underbrace{A^2K_m\left(p\ln(n/K_m)+4\log(n)\right)+2a\ln(2)\log(n)+O(1)}_{\leq O(p\log^2(n))},\end{aligned}$$

where the upper bound on the regret follows when $K_m = O(\log(n))$. This proves the upper bound in Theorem 1 and concludes the construction of the algorithm. Before we conclude the proof, we finally discuss the computational complexity of the introduced algorithm to in detail.

The computational complexity for the construction of the incremental decision tree is $O(|\mathcal{P}(\kappa)|)$, where κ represents a leaf node of the incremental decision tree (see lines $2-35$ of the algorithm in Fig. 4 and note that $|\mathcal{T}_\kappa| \leq |\mathcal{P}(\kappa)|$). The computational complexity of the sequential weighting method is $O(|\mathcal{P}(\kappa)|)$ (see (6) and lines $36-49$ of the algorithm in Fig. 4). According to the incremental hierarchical partitioning method described, the number of light nodes on the tree (see Fig. 3) is t at time t, therefore we may observe a decision tree of depth n, i.e., $|\mathcal{P}(\kappa)| = n$, in the worst-case scenario, e.g., when $\boldsymbol{x}[t] = [A, \ldots, A]^T$ for all t. Hence, the computational complexity of the algorithm over a data length of n is upper bounded by $O(n)$. Although theoretically the computational complexity of the algorithm is upper bounded by $O(n)$, in many real life applications the regressor vectors converge to stationary distributions [16]. Hence, in such practical applications, the computational complexity of the algorithm can be upper bounded by $O(\log(n))$ as discussed in Remark 2. We emphasize that in order to achieve the computational complexity $O(\log(n))$, we do not require any statistical assumptions, instead it is sufficient that the regressor vectors are evenly (to some degree) distributed in the regressor space. This concludes the proof of Theorem 1. □

for $t = 1$ **to** n **do**
  % Find the set of nodes containing $\boldsymbol{x}[t]$
  $\kappa = \lambda$
  $\mathcal{S} = \kappa$
  while s has children **do**
    $\kappa = \kappa\nu$, where ν is the last letter of the child containing $\boldsymbol{x}[t]$.
    $\mathcal{S} = \mathcal{S} + \kappa$
  end while
  % Check the index of the leaf node κ: if $\alpha_\kappa = 0$, tree remains the same.
  if $\alpha_\kappa = 0$ **then**
    $\alpha_\kappa = \alpha_\kappa + 1$
    $\mathcal{T}_\kappa = \mathcal{T}_\kappa + t$
    % If $\alpha_\kappa = 1$, create nodes $\kappa 0$ and $\kappa 1$.
  else
    % Train nodes $\kappa 0$ and $\kappa 1$.
    for all $z \in \mathcal{T}_\kappa$ **do**
      if $\boldsymbol{x}[z] \in \mathcal{R}_{\kappa 0}$ **then**
        $\mathcal{T}_{\kappa 0} = \mathcal{T}_{\kappa 0} + z$
        $L_{\kappa 0} = L_{\kappa 0} \exp(-(d[z] - \boldsymbol{w}_{\kappa 0}^T \boldsymbol{x}[z])^2/2a)$
        $P_{\kappa 0} = L_{\kappa 0}$
        $\boldsymbol{R}_{\kappa 0} = \boldsymbol{R}_{\kappa 0} + \boldsymbol{x}[z]\boldsymbol{x}[z]^T$
        $\boldsymbol{w}_{\kappa 0} = \boldsymbol{w}_{\kappa 0} + \boldsymbol{R}_{\kappa 0} \backslash (\boldsymbol{x}[z](d[z] - \boldsymbol{w}_{\kappa 0}^T \boldsymbol{x}[z]))$
      else
        % Do the similar for node $\kappa 1$.
      end if
    end for
    for all $\kappa \in \mathcal{S}$ **do**
      $P_\kappa = (P_{\kappa 0} P_{\kappa 1} + L_\kappa)/2$
    end for
    % Find the child containing $\boldsymbol{x}[t]$ and perform tree updates.
    $\nu = 0$, if $\boldsymbol{x}[t] \in \mathcal{R}_{\kappa 0}$, $\nu = 1$, otherwise.
    $\kappa = \kappa\nu$
    $\mathcal{S} = \mathcal{S} + \kappa$
    $\alpha_\kappa = 1$
  end if
  % Calculate combination weights and perform estimation.
  for all $\kappa_i \in \mathcal{P}(\kappa)$ **do**
    if $\kappa_i = \lambda$ **then**
      $\pi_{\kappa_i} = 1/2$
    else if $\kappa_i \notin \{\lambda, \kappa\}$ **then**
      $\pi_{\kappa_i} = P_{\kappa_{i-1}\nu_i^c} \pi_{\kappa_{i-1}}/2$
    else
      $\pi_{\kappa_i} = P_{\kappa_{i-1}\nu_i^c} \pi_{\kappa_{i-1}}$
    end if
    $\mu_{\kappa_i} = \pi_{\kappa_i} L_{\kappa_i}/P_\lambda$
    $\hat{d}_{\kappa_i} = \boldsymbol{w}_{\kappa_i}^T \boldsymbol{x}[t]$
  end for
  $\hat{d} = \boldsymbol{\mu}^T \hat{\boldsymbol{d}}$
  $e = d[t] - \hat{d}$
  % Perform algorithmic updates.
  for all $\kappa_i \in \mathcal{P}(\kappa)$ **do**
    $L_{\kappa_i} = L_{\kappa_i} \exp(-(d[t] - \hat{d}_{\kappa_i})^2/(2a))$
    if $\kappa_i = \kappa$ **then**
      $P_{\kappa_i} = L_{\kappa_i}$
    else
      $P_{\kappa_i} = (P_{\kappa_i 0} P_{\kappa_i 1} + L_{\kappa_i})/2$
    end if
    $R_{\kappa_i} = R_{\kappa_i} + \boldsymbol{x}[t]\boldsymbol{x}[t]^T$
    $w_{\kappa_i} = w_{\kappa_i} + R_{\kappa_i} \backslash (\boldsymbol{x}[t](d[t] - \hat{d}_{\kappa_i}))$
  end for
end for

Fig. 4: The pseudocode of the Incremental Decision Tree (IDT) regressor

Remark 1: Note that the algorithm in Fig. 4 achieves the performance of the best piecewise linear model having $O(\log(n))$ partitions with a regret of $O(p \log^2(n))$. In the most generic case, i.e., for an arbitrary piecewise model m having $O(K_m)$ partitions, the introduced algorithm still achieves a regret of $O(pK_m \log(n/K_m))$. This indicates that for models having $O(n)$ partitions, the introduced algorithm achieves a regret of $O(pn)$, hence the performance of the piecewise model cannot be asymptotically achieved. However, we emphasize that no other algorithm can achieve a smaller regret than $O(pn)$ [8], i.e., the introduced algorithm is optimal in a strong minimax sense. Intuitively, this lower bound can be justified by considering the case, in which the regressor vector at time t falls into the tth region of the piecewise model.

Remark 2: As mentioned in Remark 1 (and also can be observed in (7)), no algorithm can converge to the performance of the piecewise linear models having $O(n)$ disjoint regions. Therefore, we can limit the maximum depth of the tree by $O(\log(t))$ at each time t to achieve a low complexity implementation. With this limitation and according to the update rule of the tree, we can observe that while dividing a region into two disjoint regions, we may be forced to perform $O(t)$ computations due to the accumulated regressor vectors (since we no longer have $|\mathcal{T}_\kappa| \leq |\mathcal{P}(\kappa)|$ but instead have $|\mathcal{T}_\kappa| \leq t$). However, since a regressor vector is processed by at most $O(\log(n))$ nodes for any n, the average computational complexity of the update rule of the tree remains $O(\log(n))$. Furthermore, the performance of this low complexity implementation will be asymptotically the same as the exact implementation provided that the regressor vectors are evenly distributed in the regressor space, i.e., they are not gathered around a considerably small neighborhood. This result follows when we multiply the tree construction regret in (2) by the total number of accumulated regressor vectors, whose order, according to the above condition, is upper bounded by $o(n/\log(n))$.

Remark 3: We emphasize that the node indexes, i.e., α_κ's, determines when to create finer regions. According to the described procedure, if a node at depth l is partitioned into smaller regions, then its ith predecessor, i.e., $\kappa_i \in \mathcal{P}(\kappa)$, has observed at least $l-i$ different regressor vectors. Hence, a child node is created when coarser regions (i.e., predecessor nodes)

are sufficiently trained. In this sense, we introduce new nodes to the tree according to the current status of the tree as well as the most recent data. We also point out that, in this paper, we divide each region from its midpoint (see Fig. 3) to maintain universality. However, this process can also be performed in a data dependant manner, e.g., one can partition each region using the hyperplane that is perpendicular to the line joining two regressor vectors in that region. If there are more than two accumulated regressor vectors, then more advanced methods such as support vectors and anomaly detectors can be used to define a separator hyperplane. All these methods can be straightforwardly incorporated into our framework to produce different algorithms depending on the regression task.

D. Proof of Theorem 2

We begin our proof by emphasizing that the introduced algorithm converges to the best linear model in each region with a regret of $O(p\log^2(n))$ for any finite regression parameter $\boldsymbol{v}_m$ (since $||\boldsymbol{v}_m|| \leq \delta G p \log(n)$) as already proven in Theorem 1. Therefore, using any other linear model yields a higher regret. Hence, say we define a suboptimal affine model by applying Taylor's theorem to a twice differentiable function $f \in \mathcal{F}$ about the midpoint of each region. Let $\hat{d}_s[t]$ denote the prediction of this suboptimal affine regressor. Then, we have

$$\sum_{t=1}^{n}\left(d[t]-\hat{d}[t]\right)^2 \leq \sum_{t=1}^{n}\left(d[t]-\hat{d}_s[t]\right)^2 + O(p\log^2(n)).$$

Now applying the mean value theorem with the Lagrange form of the remainder, we obtain

$$\sum_{t=1}^{n}\left(d[t]-\hat{d}_s[t]\right)^2 - \sum_{t=1}^{n}\left(d[t]-\hat{d}_f[t]\right)^2$$
$$\leq 2A\sum_{t=1}^{n}\left\{\sum_{i=1}^{p}\sum_{j=1}^{p}\left.\frac{\partial^2 f(\boldsymbol{x})}{\partial x_i \partial x_j}\right|_{\boldsymbol{x}=\boldsymbol{b}}(x_i[t]-a_{\kappa,i})(x_j[t]-a_{\kappa,j})\right\},$$

for some $m \in \mathcal{M}_n'$ and $\boldsymbol{b} \in \mathcal{R}_\kappa$, where $\boldsymbol{a}_\kappa \triangleq [a_{\kappa,1},\ldots,a_{\kappa,p}]^T$ is the midpoint of the region $\mathcal{R}_\kappa$. Maximizing this upper bound with respect to $\boldsymbol{x}$ we obtain

$$\sum_{t=1}^{n}\left(d[t]-\hat{d}[t]\right)^2 - \sum_{t=1}^{n}\left(d[t]-\hat{d}_f[t]\right)^2$$
$$\leq 2ADp^2 n\frac{A^2}{O(\log^{2/p}(n))} + O(p\log^2(n))$$
$$\leq o(p^2 n).$$

This concludes the proof of Theorem 2. □

V. Simulations

In this section, we investigate the performance of the introduced algorithm with respect to various methods under several benchmark scenarios. Throughout the experiments, we denote the incremental decision tree algorithm of Theorem 1 by "IDT", the context tree weighting algorithm of [8] by "CTW", the linear regressor by "LR", the Volterra series regressor by "VSR" [6], the sliding window Multivariate Adaptive Regression Splines of [31], [32] by "MARS", and the Fourier nonlinear regressor of [19] by "FNR". The combination weights of the LR, VSR, and FNR are updated using the recursive least squares (RLS) algorithm [16]. Unless otherwise stated, the CTW algorithm has depth 2, the VSR, FNR, and MARS algorithms are second order, and the MARS algorithm uses 21 knots with a window length of 500 that shifts in every 200 samples.

Algorithm	Computational Complexity
IDT	$O\left(p^2\log(n)\right)$
CTW	$O\left(p^2 d\right)$
LR	$O\left(p^2\right)$
VSR	$O\left(p^{2r}\right)$
MARS	$O\left(rbw^3\right)$
FNR	$O\left((pr)^{2r}\right)$

TABLE I: Comparison of the computational complexities of the proposed algorithms with the corresponding update rules. In the table, p represents the dimensionality of the regressor space, d represents the depth of the trees in the respective algorithms, and r represents the order of the corresponding filters and algorithms. For the MARS algorithm (particularly, the fast MARS algorithm, cf. [32]), b represents the number of basis functions and w represents the window length.

In Table I, we provide the computational complexities of the proposed algorithms. We emphasize that although the computational complexity to create and run the incremental decision tree is $O(\log(n))$, the overall computational complexity of the algorithm is $O(p^2\log(n))$ due to the universal linear regressors at each region. Particularly, since the universal linear regressor at each region has a computational complexity of $O(p^2)$, the overall computational complexity of $O(p^2\log(n))$ follows. However, this universal linear regressor can be straightforwardly replaced with any linear (or nonlinear) regressor in the literature. For example, if we use the LMS algorithm to update the parameters of the linear regressor instead of using the universal algorithm for this update, the computational complexity of the overall structure becomes $O(p\log(n))$. Hence, although the computational complexity of the original IDT algorithm is $O(\log(n))$, this computational complexity may increase according to the computational complexity of the node regressors.

In this section, we first illustrate the performances of the proposed algorithms for a synthetic piecewise linear model that do not match the modeling structure of any of the above algorithms. We then consider the prediction of chaotic signals (generated from Duffing and Tinkerbell maps) and well-known data sequences such as Mackey-Glass sequence and Chua's circuit [7]. Finally, we consider the prediction of real life examples that can be found in various benchmark data set repositories such as [33], [34].

A. Synthetic Data

In this subsection, we consider the scenario where the desired data is generated by the following piecewise linear model

$$d[t] = \begin{cases} x_1[t]+x_2[t]+n[t] & \text{, if } ||\boldsymbol{x}[t]||^2 \in [0,0.1]\cup[0.5,1] \\ -x_1[t]-x_2[t]+n[t] & \text{, otherwise} \end{cases}, \quad (8)$$

and $\boldsymbol{x}[t] = [x_1[t], x_2[t]]^T$ are sample functions of a jointly Gaussian process of mean $[0,0]^T$ and covariance matrix $\boldsymbol{I}$, and

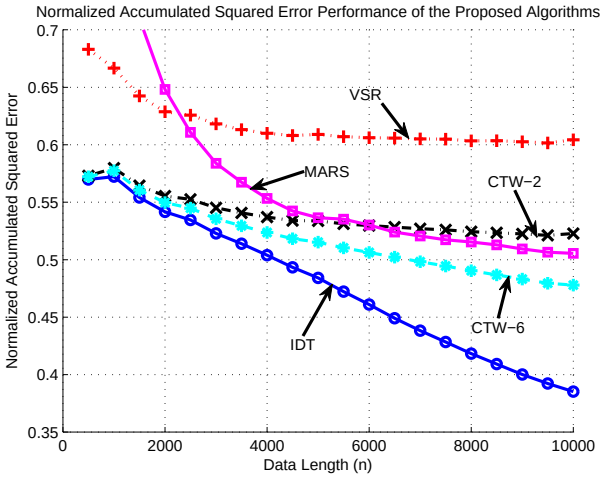

Fig. 5: Normalized accumulated squared error performances for the piecewise linear model in (8) averaged over 10 trials.

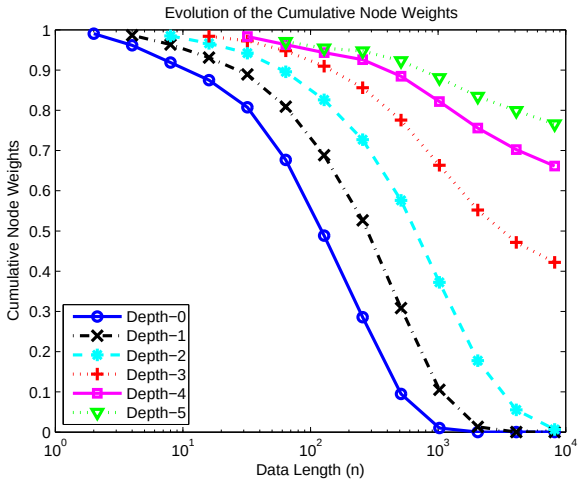

Fig. 6: Evolution of the normalized cumulative node weights at the corresponding depths of the tree for the piecewise linear model in (8) averaged over 10 trials.

$n[t]$ is a sample function from a zero mean white Gaussian process with variance 0.1. Note that the piecewise model in (8) has circular regions, which cannot be represented by hyperplanes or twice differentiable functions. Hence, the underlying relationship between the desired data and the regressor vectors cannot be exactly modeled using any of the proposed algorithms.

In Fig. 5, we present the normalized accumulated squared errors of the proposed algorithms averaged over 10 trials. For this experiment, "CTW-2" and "CTW-6" show the performances of the CTW algorithm with depths 2 and 6, respectively. Since the performances of the LR and FNR algorithms are incomparable with the rest of the algorithms, they are not included in the figure for this experiment. Fig. 5 illustrates that even for a highly nonlinear system (8), our algorithm significantly outperforms the other algorithms. The normalized accumulated error of the introduced algorithm goes to the variance of the noise signal as n increases, unlike the rest of the algorithms, whose performances converge to the performance of their optimal batch variants as n increases. This observation can be seen in Fig. 5, where the normalized cumulative error of the IDT algorithm steadily decreases since the IDT algorithm creates finer regions as the observed data length increases. Hence, even for a highly nonlinear model such as the circular piecewise linear model in (8), which cannot be represented via hyperplanes, the IDT algorithm can well approximate this highly nonlinear relationship by incrementally introducing finer partitions as the observed data length increases.

Furthermore, even though the depth of the introduced algorithm is comparable with the CTW-6 algorithm over short data sequences, the performance of our algorithm is superior with respect to the CTW-6 algorithm. This results since the IDT algorithm intrinsically eliminates the extremely finer models at the early processing stages and introduces them whenever they are needed, unlike the CTW-6 algorithm. This procedure can be observed in Fig. 6, where the IDT algorithm introduces finer regions (i.e., nodes with higher depths) to the hierarchical model as the coarser regions becomes unsatisfactory. Since the universal algorithms such as CTW distribute a "budget" into numerous experts, as the number of experts increases, the performance of such algorithms deteriorate. On the other hand, the introduced algorithm intrinsically limits the number of experts according to the unknown data length at each iteration, hence we avoid such possible performance degradations as can be observed in Fig. 6.

B. Chaotic Data

In this subsection, we consider prediction of the chaotic signals generated from the Duffing and Tinkerbell maps. The Duffing map is generated by the following discrete time equation

$$x[t+1] = ax[t] - (x[t])^3 - bx[t-1], \tag{9}$$

where we set $a = 2.75$ and $b = 0.2$ to produce the chaotic behavior [9], [35]. The Tinkerbell map is generated by the following discrete time equations

$$x[t+1] = (x[t])^2 - (y[t])^2 + ax[t] + by[t] \tag{10}$$

$$y[t+1] = 2x[t]y[t] + cx[t] + dy[t], \tag{11}$$

where we set $a = 0.9$, $b = -0.6013$, $c = 2$, and $d = 0.5$ [8], [35]. We emphasize that these values are selected to generate the well-known chaotic behaviors of these attractors.

Fig. 7a and Fig. 7b shows the normalized accumulated squared error performances of the proposed algorithms. We emphasize that due to the chaotic nature of the signals, we observe non-uniform curves in Fig. 7. Since the conventional nonlinear and piecewise linear regression algorithms commit to a priori partitioning and/or basis functions, their performances are limited by the performances of the optimal batch regressors using these prior partitioning and/or basis functions as can be observed in Fig. 7. Hence, such prior selections result in fundamental performance limitations for these algorithms. For example, in the CTW algorithm, the partitioning of the regressor space is set before the processing

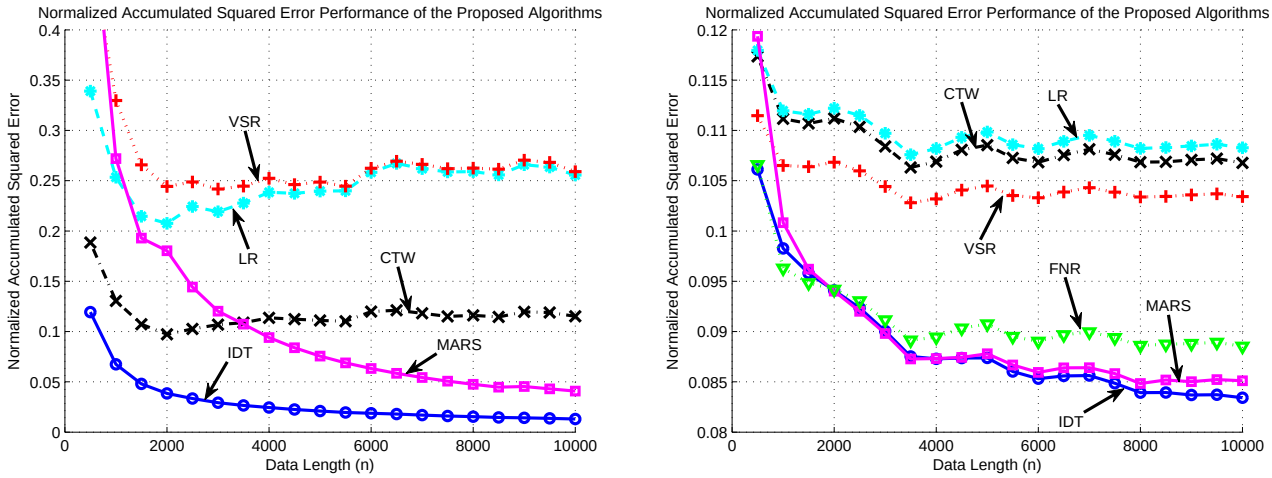

(a) **(b)**

Fig. 7: Normalized accumulated squared error performances for the chaotic data generated by (a) the Duffing map in (9) (b) the Tinkerbell map in (10) and (11).

starts. If this partitioning does not match with the underlying partitioning of the regressor space, then the performance of the CTW algorithm becomes highly unsatisfactory as seen in Fig. 7. On the other hand, the introduced algorithm illustrates a robust and superior performance while learning the underlying chaotic relationships, whereas the rest of the algorithms yields an inconsistent performance due to the chaotic nature of the signals and the limited modeling power of these algorithms.

C. Benchmark Sequences

In this subsection, we consider the prediction of the Mackey-Glass and Chua's circuit sequences. The Mackey-Glass sequence is defined by the following differential equation

$$\frac{dx[t]}{dt} = \frac{\beta x[t-\tau]}{1+(x[t-\tau])^n} - \gamma x[t], \qquad (12)$$

where we set $\beta = 2$, $\gamma = 1$, $\tau = 2$, and $n = 10$ with the initial condition $x[t] = 0.5$ for $t < 0$. We generate the time series using the fourth order Runge-Kutta method. The Chua's circuit is generated according to the following differential equations

$$\frac{dx}{dt} = \alpha(y - x - f(x)),\ \frac{dy}{dt} = x - y + z,\ \frac{dz}{dt} = -\beta y, \quad (13)$$

where we dropped the time index for notational simplicity, and $f(x) = m_1 x + 0.5(m_0 - m_1)(|x+1| - |x-1|)$, $\alpha = 15.6$, $\beta = 28$, $m_0 = -1.143$, $m_1 = -0.714$ with initial conditions $[x, y, z] = [0.7, 0, 0]$ for $t < 0$.

Fig. 8a and Fig. 8b shows the normalized accumulated squared error performances of the proposed algorithms for the Mackey-Glass and Chua's circuit sequences, respectively. We emphasize that due to the chaotic nature of the signals, we observe non-uniform curves in Fig. 8. In the figures, the algorithms with incomparable (i.e., unsatisfactory) performance are omitted. Fig. 8 presents that the IDT algorithm achieves an average of 20% relative gain in the performance with respect to the other algorithm and can accurately predict these well-known data sequences.

D. Real Data

In this subsection, we evaluate the performance of the proposed algorithms for two well-known real data sets in machine learning literature, namely "kinematics" and "pumadyn" [33], [34]. The kinematics data set involves a realistic simulation of the forward dynamics of an 8 link all-revolute robot arm and the task is to predict the distance of the end-effector from a target. Among its variants, we used the one having 9 attributes and being nonlinear as well as medium noisy. The pumadyn data set involves a realistic simulation of the dynamics of a Puma 560 robot arm. The task in these datasets is to predict the angular acceleration of one of the robot arm's links. Among its variants, we used the one having 9 attributes and being nonlinear as well as medium noisy.

Fig. 9a and Fig. 9b shows the normalized accumulated squared error performances of the proposed algorithms for the kinematics and pumadyn data sets, respectively. In the experiments, all dimensions of the regressor vector and desired data are normalized between $[-1, 1]$. Although in Fig. 9a, the VSR algorithm provides the best performance and in Fig. 9b, the MARS algorithm achieves the minimum accumulated error, the performances of these algorithms in the reciprocal experiments are highly unsatisfactory. This result implies that the data in the first experiment can be well approximated by Volterra series, whereas the model that generates the data in the second experiment is more inline with B-splines. Hence, the performances of these algorithms are extremely sensitive to the underlying structure that generates the data. On the other hand, the IDT algorithm nearly achieves the performance of the best algorithm in both experiments and presents a desirable performance under different scenarios. This result implies that the introduced algorithm can be used in various frameworks without any significant performance degradations owing to its guaranteed performance upper bounds without any statistical or structural assumptions.

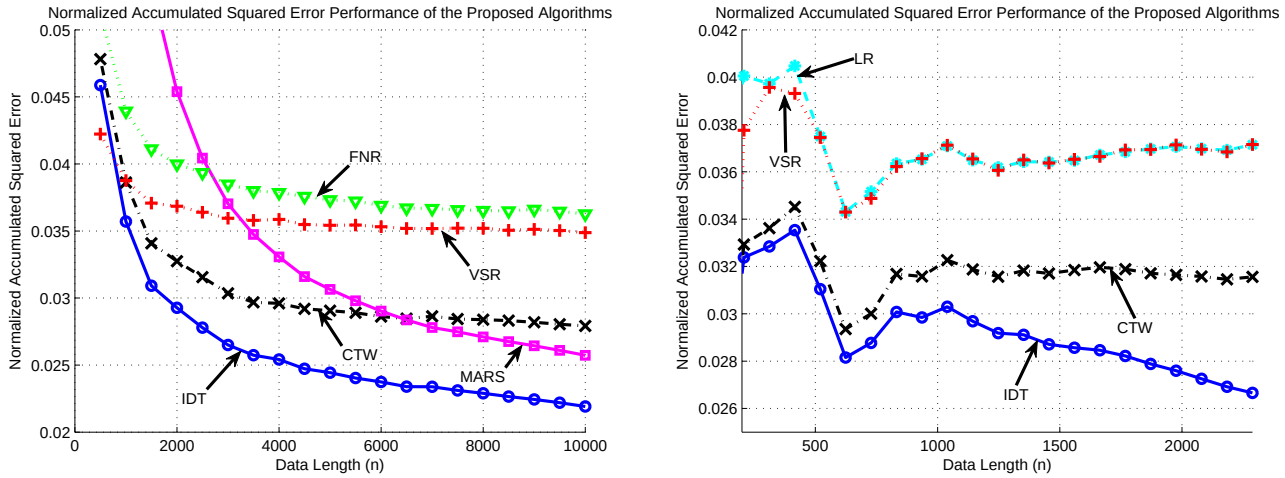

(a) **(b)**

Fig. 8: Normalized accumulated squared error performances for (a) the Mackey-Glass sequence in (12) (b) the Chua's circuit sequence in (13).

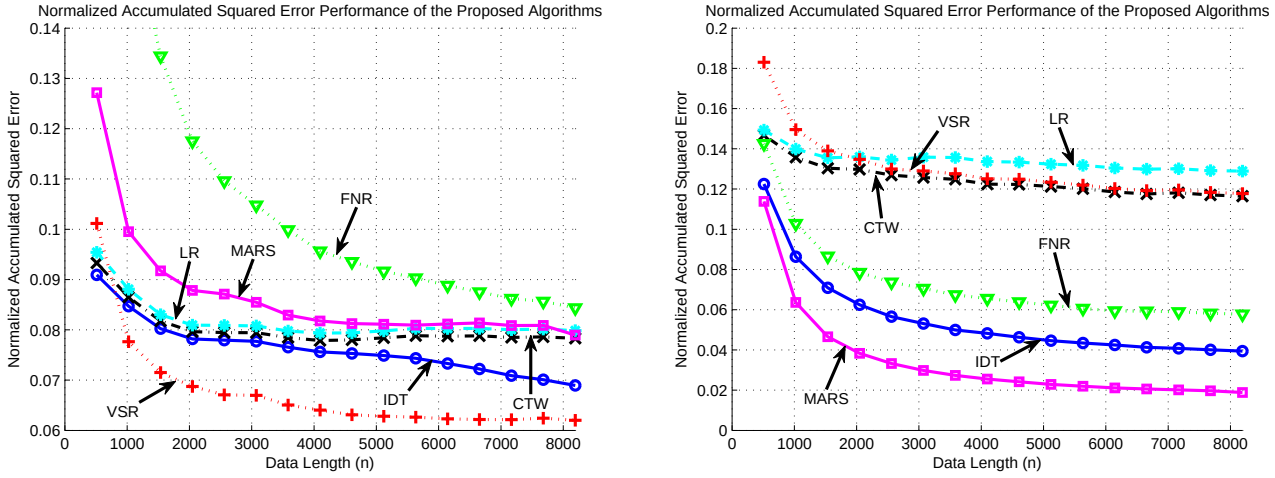

(a) **(b)**

Fig. 9: Normalized accumulated squared error performances for (a) the "kinematics" (b) the "pumadyn" data sets.

VI. Concluding Remarks

We study nonlinear regression of deterministic signals using an incremental decision tree, where the regressor space is partitioned using a nested structure and independent regressors are assigned to each region. In this framework, we introduce a tree based algorithm that sequentially increases its nonlinear modeling power and achieves the performance of the optimal twice differentiable function as well as the performance of the best piecewise linear model defined on the incremental decision tree. Furthermore, this performance is achieved only with a computational complexity logarithmic in the data length n, i.e., $O(\log(n))$ (under regularity conditions). We demonstrate the superior performance of the introduced algorithm over a series of well-known benchmark applications in the regression literature.

References

[1] L. Devroye, T. Linder, and G. Lugosi, "Nonparametric estimation and classification using radial basis function nets and empirical risk minimization," *IEEE Transactions on Neural Networks*, vol. 7, no. 2, pp. 475–487, Mar 1996.

[2] A. Krzyzak and T. Linder, "Radial basis function networks and complexity regularization in function learning," *IEEE Transactions on Neural Networks*, vol. 9, no. 2, pp. 247–256, Mar 1998.

[3] N. D. Vanli and S. S. Kozat, "A unified approach to universal prediction: Generalized upper and lower bounds," *IEEE Transactions on Neural Networks and Learning Systems*, vol. PP, no. 99, pp. 1–1, 2014.

[4] R. Gribonval, "From projection pursuit and CART to adaptive discriminant analysis?" *IEEE Transactions on Neural Networks*, vol. 16, no. 3, pp. 522–532, May 2005.

[5] I. Ali and Y.-T. Chen, "Design quality and robustness with neural networks," *IEEE Transactions on Neural Networks*, vol. 10, no. 6, pp. 1518–1527, Nov 1999.

[6] M. Schetzen, *The Volterra and Wiener Theories of Nonlinear Systems*. NJ: John Wiley & Sons, 1980.

[7] O. J. J. Michel, A. O. Hero, and A.-E. Badel, "Tree-structured nonlinear signal modeling and prediction," *IEEE Transactions on Signal Processing*, vol. 47, no. 11, pp. 3027–3041, 1999.

[8] S. S. Kozat, A. C. Singer, and G. C. Zeitler, "Universal piecewise linear prediction via context trees," *IEEE Transactions on Signal Processing*, vol. 55, no. 7, pp. 3730–3745, 2007.

[9] Y. Yilmaz and S. S. Kozat, "Competitive randomized nonlinear prediction under additive noise," *IEEE Signal Processing Letters*, vol. 17, no. 4, pp. 335–339, April 2010.

[10] Q. Zhu, Z. Zhang, Z. Song, Y. Xie, and L. Wang, "A novel nonlinear regression approach for efficient and accurate image matting," *IEEE Signal Processing Letters*, vol. 20, no. 11, pp. 1078–1081, 2013.

[11] H. Zhao and J. Zhang, "A novel adaptive nonlinear filter-based pipelined feedforward second-order volterra architecture," *IEEE Transactions on Signal Processing*, vol. 57, no. 1, pp. 237–246, 2009.

[12] D. P. Helmbold and R. E. Schapire, "Predicting nearly as well as the best pruning of a decision tree," *Machine Learning*, vol. 27, no. 1, pp. 51–68, 1997.

[13] R. J. Drost and A. C. Singer, "Constrained complexity generalized context-tree algorithms," in *IEEE/SP 14th Workshop on Statistical Signal Processing*, 2007, pp. 131–135.

[14] E. Takimoto, A. Maruoka, and V. Vovk, "Predicting nearly as well as the best pruning of a decision tree through dyanamic programming scheme," *Theoretical Computer Science*, vol. 261, pp. 179–209, 2001.

[15] E. Takimoto and M. K. Warmuth, "Predicting nearly as well as the best pruning of a planar decision graph," *Theoretical Computer Science*, vol. 288, pp. 217–235, 2002.

[16] A. H. Sayed, *Fundamentals of Adaptive Filtering*. NJ: John Wiley & Sons, 2003.

[17] V. Kekatos and G. Giannakis, "Sparse Volterra and polynomial regression models: Recoverability and estimation," *IEEE Transactions on Signal Processing*, vol. 59, no. 12, pp. 5907–5920, 2011.

[18] M. Scarpiniti, D. Comminiello, R. Parisi, and A. Uncini, "Nonlinear spline adaptive filtering," *Signal Processing*, vol. 93, no. 4, pp. 772 – 783, 2013.

[19] A. Carini and G. L. Sicuranza, "Fourier nonlinear filters," *Signal Processing*, vol. 94, no. 0, pp. 183 – 194, 2014.

[20] N. D. Vanli and S. S. Kozat, "A comprehensive approach to universal piecewise nonlinear regression based on trees," *IEEE Transactions on Signal Processing*, vol. 62, no. 20, pp. 5471–5486, Oct 2014.

[21] S. Dasgupta and Y. Freund, "Random projection trees for vector quantization," *IEEE Transactions on Information Theory*, vol. 55, no. 7, pp. 3229–3242, 2009.

[22] J. L. Bentley, "Multidimensional binary search trees in database applications," *IEEE Transactions on Software Engineering*, vol. SE-5, no. 4, pp. 333–340, 1979.

[23] N. B. Lee, A. B. and L. Wasserman, "Treelets – an adaptive multiscale basis for sparse unordered data," *The Annals of Applied Statistics*, vol. 2, no. 2, pp. 435–471, 2008.

[24] J. Gama, "Functional trees," *Machine Learning*, vol. 55, no. 3, pp. 219–250, 2004.

[25] A. V. Aho and N. J. A. Sloane, "Some doubly exponential sequences," *Fibonacci Quarterly*, vol. 11, pp. 429–437, 1970.

[26] F. M. J. Willems, Y. M. Shtarkov, and T. J. Tjalkens, "The context-tree weighting method: basic properties," *IEEE Transactions on Information Theory*, vol. 41, no. 3, pp. 653–664, 1995.

[27] S. S. Kozat, A. T. Erdogan, A. C. Singer, and A. H. Sayed, "Steady state MSE performance analysis of mixture approaches to adaptive filtering," *IEEE Transactions on Signal Processing*, vol. 58, no. 8, pp. 4050–4063, August 2010.

[28] J. Arenas-Garcia, A. R. Figueiras-Vidal, and A. H. Sayed, "Mean-square performance of a convex combination of two adaptive filters," *IEEE Transactions on Signal Processing*, vol. 54, pp. 1078–1090, 2006.

[29] A. C. Singer, S. S. Kozat, and M. Feder, "Universal linear least squares prediction: upper and lower bounds," *IEEE Transactions on Information Theory*, vol. 48, no. 8, pp. 2354–2362, 2002.

[30] F. M. J. Willems, "Coding for a binary independent piecewise-identically-distributed source," *IEEE Transactions on Information Theory*, vol. 42, no. 6, pp. 2210–2217, Nov 1996.

[31] J. H. Friedman, "Multivariate adaptive regression splines," *The Annals of Statistics*, vol. 19, no. 1, pp. 1–67, 1991.

[32] ——, "Fast MARS," *Stanford University Technical Report*, 1993. [Online]. Available: http://www.milbo.users.sonic.net/earth/Friedman-FastMars.pdf

[33] C. Blake and C. Merz, "UCI Machine Learning Repository." [Online]. Available: http://archive.ics.uci.edu/ml/index.html

[34] C. E. Rasmussen, R. M. Neal, G. Hinton, D. Camp, M. Revow, Z. Ghahramani, R. Kustra, and R. Tibshirani, "Delve data sets." [Online]. Available: http://www.cs.toronto.edu/~delve/data/datasets.html

[35] S. Wiggins, *Introduction to Applied Nonlinear Dynamical Systems and Chaos*. Springer New York, 2003, vol. 2.